arXiv:1610.08244v2 [stat.ME] 29 Dec 2016

An Efficient Estimation of Conway-Maxwell Poisson Regression and Additive Model with an Application to Bike Sharing

Suneel Babu Chatla *
and
Galit Shmueli
Institute of Service Science
National Tsing Hua University

Abstract

The Conway-Maxwell Poisson (CMP) regression is a popular model for count data due to its ability to capture both under dispersion and over dispersion. However, CMP regression is limited when dealing with complex nonlinear relationships. With today's wide availability of count data, especially due to the growing collection of data on human and social behavior, there is need for count data models that can capture complex nonlinear relationships. One useful approach is additive models. However, there has been no additive model implementation for CMP. To fill this void, we first propose a flexible estimation framework for CMP regression based on iterative reweighed least squares (IRLS) and then extend this model to allow for additive components using smoothing splines. Because CMP belongs to the exponential family, convergence is guaranteed for IRLS and it is more efficient than the numerical optimization. We illustrate the usefulness of this approach through extensive simulation studies and using real data from a bike sharing system in Washington, DC.

Keywords: IRLS, Smoothing Splines, Backfitting , Over and under dispersion, Time series

*The authors gratefully acknowledge Prof. Li-Shan Huang and Prof. Satish Iyengar for their valuable feedback and suggestions.

1 Introduction

1.1 Motivation

Count data have become popular outcome variables in studies in various areas, especially due to the growing availability of data on human and social behavior. One example is bicycle sharing programs, which have received increasing attention in recent years and have spread in many cities around the world. Growing concerns about global climate change, energy security and excessive motorization have led the widespread adoption of sustainable transportation strategies (Shaheen et al. 2010). The principle of bike sharing is simple: Individuals use bicycles whenever needed without the cost and responsibilities of bike ownership. Bike sharing programs typically cover bike purchase, maintenance, parking, and storage responsibilities. Therefore, there is a need for bike sharing programs to effectively understand the factors that influence demand so that they can better maintain inventory, schedule repairs, and manage resources.

Many external factors influence bike demand including weather conditions, holidays, and working hours. Modeling bike sharing data therefore requires modeling time series of count data. The most commonly used regression models for cross-sectional count data are Poisson regression and Negative-Binomial regression. In addition, the Conway-Maxwell Poisson distribution (CMP) has gained increasing popularity for its flexibility and ability to handle both over- and under-dispersed data. Revived by Shmueli et al. (2005), CMP is a two-parameter generalization of the Poisson, Bernoulli, and Geometric distributions. Suppose Y is a random variable that follows a CMP distribution, then the probability mass function (p.m.f) for $Y \geq 0$ is defined as

$$P(Y = y) = \frac{\lambda^y}{(y!)^\nu \zeta(\lambda, \nu)},\ y \in \mathbb{Z}^+$$

with

$$\zeta(\lambda, \nu) = \sum_{s=0}^{\infty} \frac{\lambda^s}{(s!)^\nu}$$

for the parameters $(\lambda, \nu) \in \Theta \subseteq (0, \infty) \times (0, \infty)$.

The CMP distribution includes three well-known distributions as special cases: Poisson ($\nu = 1$), Geometric ($\nu = 0, \lambda < 1$), and Bernoulli ($\nu \to \infty$ with probability $\frac{\lambda}{1+\lambda}$). Due to the additional parameter ν, the CMP distribution is flexible enough to handle both over-dispersion ($\nu < 1$) and under-dispersion ($\nu > 1$) which are common in count data (Sellers & Shmueli 2010).

One of the major limitations of the CMP distribution is that it does not belong to the natural parameter exponential family so we cannot model the mean directly as a function of covariates, as in standard models such as Poisson and logistic regression. However, the CMP distribution belongs to the canonical parameter exponential family and thus it retains all the properties and advantages of that family.

Unfortunately, these cross sectional models are often too restrictive for modeling non-linear relationships or time series data. At the same time, additive models are widely used for modeling nonlinear relationships such as time series (Dominici et al. 2002, Stieb et al. 2003). Additive models have the advantage of being parsimonious while at the same time providing more flexibility to capture complicated relationships. However, there exist no additive model implementation for the CMP model. Existing additive model implementations are heavily dependent upon the IRLS estimation framework, which currently does not exist for CMP. In this study, we propose and implement an IRLS estimation framework for CMP and then extend that to CMP additive models.

Motivated by the need for flexible count data regression models for applications such as bike sharing, which can assist service providers in better management of their resources,

we develop a CMP additive model. The outline of this paper is follows. In Section 2 we describe the background of bike sharing data, introduce a specific bike sharing dataset, and describe the CMP regression and the problems associated with IRLS implementation. In Section 3 we develop an IRLS framework for estimating a CMP model and evaluate our proposed approach with the existing methods using an extensive simulation study. In Section 5 we use the IRLS framework to develop an additive model for CMP, and again evaluate its performance using simulation. We return to the bike sharing data and case in Section 6 and use our proposed additive model to draw valuable insights from the data. Section 7 presents conclusions and future directions.

2 Background for Data and Method

2.1 Bike sharing data

Bike sharing systems can be thought of as a new generation of traditional bike rental services where the entire process that includes membership, bike rental and return has become automated. Through these systems, users are able to easily rent a bike from a particular station and return it to another location. Data collected by bike sharing systems typically include information on each trip taken (time stamps and locations of rental and return) and sometimes also information on the rider. Several datasets from real bike sharing systems are publicly available. We use the data available from Fanaee-T & Gama (2013) on rides in Washington, DC between 2011-2012. The data is available in two formats: daily and hourly number of rentals. We chose the hourly data as it is more complex and better illustrates the new models that we introduce.

The data includes information about the number of rides by casual users and regis-

tered users for every hour during the years 2011 and 2012. In addition, it also includes external information such whether the hour is on a weekday, a working day or a holiday, the weather situation (clear/cloudy/rainy), temperature, and wind speed. As mentioned earlier, these external factors are considered detrimental to the demand for bikes. The full list of attributes and their descriptions is given in Table 1.

Figure 1 shows rental patterns for both casual and registered users for the months of January and April in year 2012. These two months are from two different seasons. As expected, there is a daily seasonality and with some high-demand days and some low-demand days. Although there appears to be a trend over time, it is difficult to ascertain as such from the time plot.Our objective is to create a model that helps understand the impact of external factors on the number of rentals for both casual users and registered users. This is different from a forecasting goal (forecasting future demand), where existing Poisson additive models might be sufficiently accurate.

Next, we provide the background for CMP Regression and explain why there exists no IRLS estimation framework for CMP. We then describe our approach for an IRLS CMP and use it to develop an additive model. In Section 6 we return to the bike sharing example and explain the findings using our proposed method.

2.2 CMP Regression

Assume that we have a random sample of n observations $\{y_i, \mathbf{x_i^T}, \mathbf{z_i^T}\}_{i=1}^n$, where $\mathbf{x_i^T} = [1, x_{i1}, \cdots, x_{ip}]$ and $\mathbf{z_i^T} = [1, z_{i1}, \cdots, z_{iq}]$. In matrix notation, let $Y = [y_1, \ldots, y_n]^T$, $X = [\mathbf{x_1}, \ldots, \mathbf{x_n}]^T$ and $Z = [\mathbf{z_1}, \ldots, \mathbf{z_n}]^T$ with the parameter vectors $\boldsymbol{\lambda}$, $\boldsymbol{\nu}$ of size n. Then the

Name	Description
dteday	date
season	season (1:spring, 2:summer, 3:fall, 4:winter)
yr	year (0: 2011, 1:2012)
mnth	month (1 to 12)
hr	hour (0 to 23)
holiday	weather the day is holiday or not (extracted from)
weekday	day of the week
workingday	if day is neither weekend nor holiday is 1, otherwise is 0.
weathersit	1= Clear, Few clouds, Partly cloudy
	2= Mist + Cloudy, Mist + Broken clouds, Mist + Few clouds, Mist
	3= Light Snow, Light Rain + Thunderstorm + Scattered clouds
	4= Heavy Rain + Ice Pallets + Thunderstorm + Mist, Snow + Fog
temp	Normalized temperature in Celsius. The values are divided to 41 (max)
atemp	Normalized feeling temperature in Celsius. The values are divided to 50 (max)
hum	Normalized humidity. The values are divided to 100 (max)
windspeed	Normalized wind speed. The values are divided to 67 (max)
casual	count of casual users
registered	count of registered users
cnt	count of total rental bikes including both casual and registered

Table 1: Full list of attributes and their description for the bike sharing data

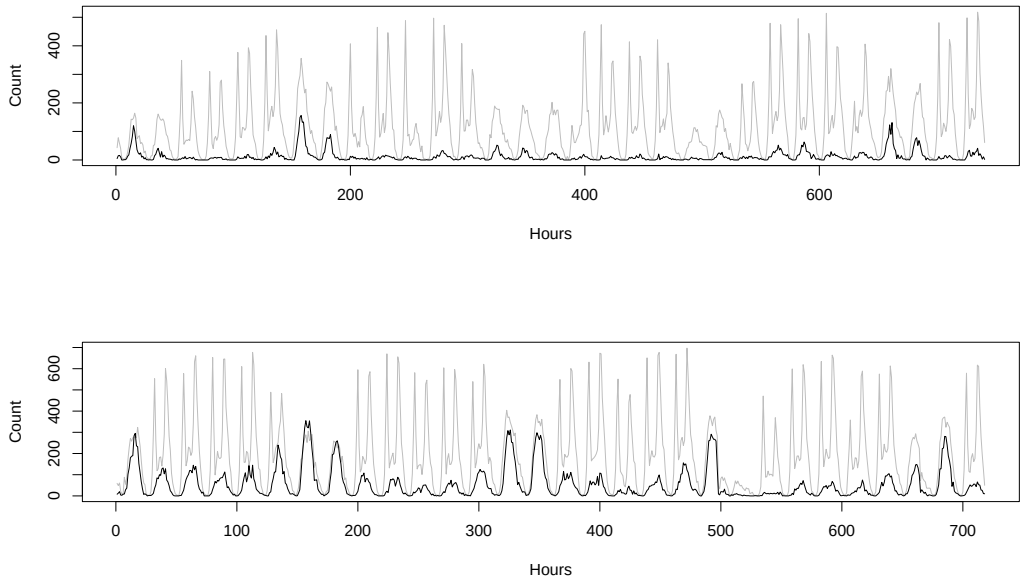

Figure 1: Number of hourly bike rentals for casual users (black) and registered users (gray) during January 2012 (top) and April 2012 (bottom)

CMP regression can be formulated as

$$\ln(\boldsymbol{\lambda}) = X\boldsymbol{\beta} \tag{1}$$

$$\ln(\boldsymbol{\nu}) = Z\boldsymbol{\gamma} \tag{2}$$

where $\boldsymbol{\beta} \in \mathbb{R}^{p+1}, \boldsymbol{\nu} \in \mathbb{R}^{q+1}$.

Using this model formulation, the log likelihood can be written as

$$\ell(Y, \boldsymbol{\beta}, \boldsymbol{\gamma}) = Y^T X\boldsymbol{\beta} - \ln(Y!)^T \exp\{Z\boldsymbol{\gamma}\} - \mathbf{1}^T \ln\zeta(\exp\{X\boldsymbol{\beta}\}, \exp\{Z\boldsymbol{\gamma}\}),$$

which yields the following score equations:

$$\frac{\partial \ell}{\partial \boldsymbol{\beta}^{\boldsymbol{T}}} = X^T \left(Y - \frac{\partial \ln \zeta}{\partial \ln \boldsymbol{\lambda}} \right) = X^T (Y - \mathbb{E}[Y])$$

$$\frac{\partial \ell}{\partial \boldsymbol{\gamma}^{\boldsymbol{T}}} = Z^T \left[\left(-\ln(Y!) + \frac{\partial \ln \zeta}{-\partial \boldsymbol{\nu}} \right) * \boldsymbol{\nu} \right] = Z^T \left[(-\ln(Y!) + \mathbb{E}[\ln(Y!)]) * \boldsymbol{\nu} \right]$$

where * denotes element-wise multiplication and $\mathbb{E}[\cdot]$ is the mean function.

Since the derivatives of ζ do not have closed form solutions, the score equations cannot be solved as in standard GLMs such as Poisson. For this reason, the existing implementations of CMP models either use numerical gradient-based methods or MCMC, but do not use iterative re-weighted least squares (IRLS), which is the workhorse routine for estimation of all the standard GLMs. Although gradient-based methods have faster convergence rate that IRLS, they are not efficient because they use the observed information matrix, and are not robust to outliers. In contrast, the IRLS is more efficient and robust but it is relatively slower than gradient-based methods(Green 1984).

Another advantage of the IRLS method is that modeling extensions such as additive models and lasso can be implemented easily (Yee 2007). To the best of our knowledge, there is no implementation of an IRLS algorithm for CMP regression. While Sellers & Shmueli (2010) did not implement an IRLS for CMP, they briefly outlined the IRLS algorithm for

CMP. Their approach is based on iteratively solving the following weighted least squares (WLS) equation:

$$\begin{bmatrix} X^T \\ (g(Y) * Z)^T \end{bmatrix} W \begin{bmatrix} X & (g(Y) * Z) \end{bmatrix} \begin{bmatrix} \boldsymbol{\beta^m} \\ \boldsymbol{\gamma^m} \end{bmatrix} = \begin{bmatrix} X^T \\ (g(Y) * Z)^T \end{bmatrix} W \quad T$$

where $W = diag(\mathbb{V}(y_1), \ldots, \mathbb{V}(y_n))$,$\mathbb{V}(\cdot)$ is the variance function, $g(Y) = \frac{-\ln(Y!)+\mathbb{E}[\ln(Y!)]}{Y-\mathbb{E}[Y]} * \boldsymbol{\nu}$ and

$$T = X\boldsymbol{\beta^{m-1}} + (g(Y) * Z)\boldsymbol{\gamma^{m-1}} + \frac{Y - \mathbb{E}[Y]}{\mathbb{V}[Y]}.$$

Originally, in their approach, the authors used $Z = [1, \ldots, 1]^T$ and did not use a log link function for $\boldsymbol{\nu}$. Here we generalize their approach using both a log link function for $\boldsymbol{\nu}$ and allowing predictors in the $\ln(\boldsymbol{\nu})$ model. While this approach looks reasonable, it has the following two drawbacks:

1. WLS does not use the expected information. For example, based on the WLS formulation the information for the intercept in $\ln(\boldsymbol{\nu})$ is

   $$\sum_{i=1}^{n} g(y_i)^2 \mathbb{V}(y_i) = \sum_{i=1}^{n} \left[\frac{-\ln(y_i!) + \mathbb{E}[\ln(y_i!)]}{y_i - \mathbb{E}[y_i]}\right]^2 \nu_i^2 \mathbb{V}[y_i] \neq \sum_{i=1}^{n} \mathbb{V}[\ln(y_i!)]\nu_i^2.$$

   The value in the right hand side is the expected information for the same $\ln(\boldsymbol{\nu})$ using score equations. Both these values can be equivalent asymptotically if we use $W* = W/n$, but not with the Sellers & Shmueli (2010) formulation.

2. The idea of combining both models into a single WLS framework is computationally attractive. However, since both $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ are dependent on each other, updating both of them in single model is problematic especially with least squares. When we implemented this approach, most of the time the algorithm remained closer to the

initial values and sometimes it chose very small values of $\boldsymbol{\nu}$ irrespective of the actual $\boldsymbol{\nu}$ value.

To overcome these limitations, we propose a two step IRLS algorithm that uses the expected information matrix for updates, with guaranteed convergence. Our approach also makes it easier to extend the CMP model for the estimation of additive components.

3 IRLS Framework for CMP

To implement the IRLS method, we must first calculate the values $\mathbb{E}[Y], \mathbb{E}[\ln(Y!)], \mathbb{V}[Y]$ and $\mathbb{V}[\ln(Y!)]$. We do this by calculating the cumulants.

3.1 Calculation of Cumulants

The standard way of calculating cumulants is to use the pmf. Since the pmf for the CMP distribution involves an infinite series, a simple approach is to truncate the infinite series in such a way that the error is bounded ($\epsilon = 10^{-6}$) Shmueli et al. (2005).

Another way of calculating the cumulants is by using the properties of the canonical parameter of the exponential family. The t^{th} cumulants for Y and $log(Y!)$ can be obtained as:

$$\kappa_t[Y] = \frac{\partial^{(t)} \ln \zeta(\boldsymbol{\lambda}, \boldsymbol{\nu})}{\partial^{(t)} \ln \boldsymbol{\lambda}}$$
$$\kappa_t[\ln(Y!)] = -\frac{\partial^{(t)} \ln \zeta(\boldsymbol{\lambda}, \boldsymbol{\nu})}{\partial^{(t)} \boldsymbol{\nu}}.$$

There has been some research trying to approximate the Z function using a closed form expression. First, Shmueli et al. (2005) provided the following asymptotic approximation:

$$\zeta(\boldsymbol{\lambda}, \boldsymbol{\nu}) = \frac{e^{\boldsymbol{\nu}\boldsymbol{\lambda}^{1/\boldsymbol{\nu}}}}{\boldsymbol{\lambda}^{\frac{\boldsymbol{\nu}-1}{2\boldsymbol{\nu}}}(2\pi)^{\frac{\boldsymbol{\nu}-1}{2}}\sqrt{\boldsymbol{\nu}}}(1 + O(\boldsymbol{\lambda}^{\frac{-1}{\boldsymbol{\nu}}})). \tag{3}$$

For higher order cumulants (order greater than 1) the cumulant generating function (CGF) has the following form:

$$\kappa_t = \boldsymbol{\nu}\boldsymbol{\lambda}^{1/\boldsymbol{\nu}}(e^{t/\boldsymbol{\nu}} - 1)$$

Although this approximation is appealing theoretically it has limited practical value. To get a better approximation with the function we should have larger $\boldsymbol{\lambda}^{1/\boldsymbol{\nu}}$ values, i.e., larger counts.

Recently, Simsek & Iyengar (2016), Gaunt et al. (2016) further improved this approximation by adding two more terms:

$$\zeta(\boldsymbol{\lambda}, \boldsymbol{\nu}) = \frac{e^{\boldsymbol{\nu}\boldsymbol{\lambda}^{1/\boldsymbol{\nu}}}}{\boldsymbol{\lambda}^{\frac{\boldsymbol{\nu}-1}{2\boldsymbol{\nu}}}(2\pi)^{\frac{\boldsymbol{\nu}-1}{2}}\sqrt{\boldsymbol{\nu}}}\left[1 + (\boldsymbol{\nu}-1)\big(A+B\big) + O(\boldsymbol{\lambda}^{\frac{-3}{\boldsymbol{\nu}}})\right], as \quad \boldsymbol{\lambda}^{1/\boldsymbol{\nu}} \to \infty, \tag{4}$$

where $A = \frac{8\boldsymbol{\nu}^2+12\boldsymbol{\nu}+3}{96\boldsymbol{\lambda}^{1/\boldsymbol{\nu}}\boldsymbol{\nu}^2}, B = \frac{1+6\boldsymbol{\nu}}{144\boldsymbol{\lambda}^{2/\boldsymbol{\nu}}\boldsymbol{\nu}^3}$.

We use the formula in (4) and calculate the required cumulants. Define $\boldsymbol{\alpha} = \boldsymbol{\lambda}^{1/\boldsymbol{\nu}}$, and $D = \mathbf{1} + (\boldsymbol{\nu}-1)(A+B)$. The cumulants are defined as follows:

$$\mathbb{E}[Y] = \boldsymbol{\alpha} - \frac{\boldsymbol{\nu}-1}{2\boldsymbol{\nu}} + \frac{1}{D}\left[(\boldsymbol{\nu}-1)\big(\frac{-A}{\boldsymbol{\nu}} - \frac{-2B}{\boldsymbol{\nu}}\big)\right]$$

$$\mathbb{V}[Y] = \frac{\boldsymbol{\alpha}}{\boldsymbol{\nu}} - \frac{1}{D}\left[(\boldsymbol{\nu}-1)\big(\frac{A}{\boldsymbol{\nu}^2} + \frac{4B}{\boldsymbol{\nu}^2}\big) - \frac{1}{D}\left((\boldsymbol{\nu}-1)\big(\frac{-A}{\boldsymbol{\nu}} - \frac{-2B}{\boldsymbol{\nu}}\big)\right)^2\right]$$

$$\begin{aligned}\mathbb{E}[\ln(Y!)] = -\Bigg[\boldsymbol{\alpha}(1 - \frac{\ln(\boldsymbol{\lambda})}{\boldsymbol{\nu}} - \frac{\ln(\boldsymbol{\lambda})}{2\boldsymbol{\nu}^2} - \frac{\ln(2\pi)}{2} - \frac{1}{2\boldsymbol{\nu}} + \frac{1}{D}\Bigg(\frac{24\boldsymbol{\nu}^2+8\boldsymbol{\nu}-9}{96\boldsymbol{\alpha}\boldsymbol{\nu}^2} - \\ (\boldsymbol{\nu}-1)(2\boldsymbol{\nu} - \ln(\boldsymbol{\lambda}))\frac{A}{\boldsymbol{\nu}^2} + \frac{1}{12\boldsymbol{\alpha}^2\boldsymbol{\nu}^2} - \big((\boldsymbol{\nu}-1)(3\boldsymbol{\nu} - 2\ln(\boldsymbol{\lambda})\frac{B}{\boldsymbol{\nu}^2}\Bigg)\Bigg]\end{aligned}$$

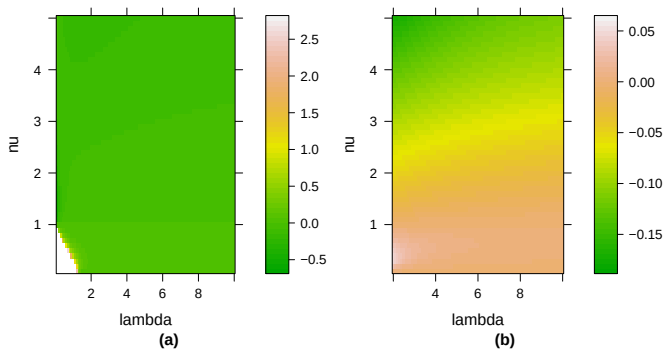

Figure 2: The differences of ζ calculated using pmf and asymptotic expressions. (a) for $0.1 \leq \lambda \leq 10$, $0.1 \leq \nu \leq 5$. (b) for $2 \leq \boldsymbol{\lambda} \leq 10$, $0.1 \leq \boldsymbol{\nu} \leq 5$.

$$\mathbb{V}[\ln(Y!)] = \frac{\ln(\boldsymbol{\lambda})}{\boldsymbol{\nu}^3}(\boldsymbol{\alpha}\ln(\boldsymbol{\lambda}) + 1) + \frac{1}{2\boldsymbol{\nu}^2} + \frac{1}{D}\left[\frac{48\boldsymbol{\nu} + 8}{96\boldsymbol{\alpha}\boldsymbol{\nu}^2} - \big(16\boldsymbol{\nu}^3 + 8\boldsymbol{\nu}^2 - 18\boldsymbol{\nu} - 6 \right.$$
$$\left.(2\boldsymbol{\nu} - \boldsymbol{\lambda})(24\boldsymbol{\nu}^2 + 8\boldsymbol{\nu} - 9)\big)\frac{1}{96\boldsymbol{\alpha}\boldsymbol{\nu}^4}\right] - (\frac{\ln(\boldsymbol{\lambda})}{\boldsymbol{\nu}} - 1)\frac{2}{12\boldsymbol{\alpha}^2\boldsymbol{\nu}^3} -$$
$$\frac{1}{144\boldsymbol{\alpha}^2\boldsymbol{\nu}^5}\Big[(12\boldsymbol{\nu} - 5)(3\boldsymbol{\nu} - 2\ln(\boldsymbol{\lambda})) + 18\boldsymbol{\nu}^2 - 15\boldsymbol{\nu} - 3\Big]$$

From Figure 2, which compares the approximation accuracy for different leels of λ and ν, it can be observed that the new approximation is reasonably good when $\boldsymbol{\lambda} \geq 2$ and $\boldsymbol{\nu} \leq 1$. For this reason we use the cumulants derived from the asymptotic expression only when $\boldsymbol{\lambda} \geq 2, \boldsymbol{\nu} \leq 1$ and for other values we use the pmf function to calculate cumulants with some bounding error. Although the approximation works for a limited range, this is very helpful because the asymptotic series converges very slowly when $\boldsymbol{\nu} < 1$ and this approximation eases the computational burden significantly.

Similarly, for the $\ln(y_i!)$ the values are computed recursively until $y_i < 254$ and after that Stirling's approximation is used as it is reasonably close (Abramowitz et al. 1966).

3.2 Two step method

Let us define $\boldsymbol{u}_{p\times 1} = \frac{\partial \ell}{\partial \boldsymbol{\beta}^{\boldsymbol{T}}}$ and $\boldsymbol{v}_{q\times 1} = \frac{\partial \ell}{\partial \boldsymbol{\gamma}^{\boldsymbol{T}}}$. The full information matrix can be written as

$$\begin{aligned} I_{p+q} &= \mathbb{E}\left[\begin{pmatrix} \boldsymbol{u} \\ \boldsymbol{v} \end{pmatrix} (\boldsymbol{u}^T \boldsymbol{v}^T) \right] \\ &= \begin{bmatrix} \mathbb{E}[\boldsymbol{u}\boldsymbol{u}^T] & \mathbb{E}[\boldsymbol{u}\boldsymbol{v}^T] \\ \mathbb{E}[\boldsymbol{v}\boldsymbol{u}^T] & \mathbb{E}[\boldsymbol{v}\boldsymbol{v}^T] \end{bmatrix} \\ &= \begin{bmatrix} I_{11} & I_{12} \\ I_{21} & I_{22} \end{bmatrix} \end{aligned}$$

with

$$\begin{aligned} \mathbb{E}[\boldsymbol{u}\boldsymbol{u}^T] &= I_{11} = X^T \Sigma_Y X \\ \mathbb{E}[\boldsymbol{u}\boldsymbol{v}^T] &= I_{12} = I_{21}^T = -\boldsymbol{\nu} * X^T \Sigma_{Y,\ln(Y!)} Z \\ \mathbb{E}[\boldsymbol{v}\boldsymbol{v}^T] &= I_{22} = \boldsymbol{\nu}^2 * Z^T \Sigma_{\ln(Y!)} Z \end{aligned}$$

where $\Sigma_Y = diag(\mathbb{V}(y_1), \ldots, \mathbb{V}(y_n)), \Sigma_{Y,\ln(Y!)} = Cov(Y, \ln(Y!))$, and
$\Sigma_{\ln(Y!)} = diag(\mathbb{V}(\ln(y_1!)), \ldots, \mathbb{V}(\ln(y_n!)))$.

Using the information matrix I, the IRLS algorithm has the following form for the m^{th} update

$$\begin{bmatrix} \boldsymbol{\beta} \\ \boldsymbol{\gamma} \end{bmatrix}^{(m)} = \begin{bmatrix} \boldsymbol{\beta} \\ \boldsymbol{\gamma} \end{bmatrix}^{(m-1)} + I^{-1} \begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{v} \end{bmatrix}$$

which implies the following equations:

$$\begin{aligned} X^T \Sigma_Y X \boldsymbol{\beta}^{(m)} - X^T \Sigma_{Y,\ln(Y!)} \boldsymbol{\nu} * Z \boldsymbol{\gamma}^{(m)} &= X^T \Sigma_Y X \boldsymbol{\beta}^{(m-1)} - \\ & \quad X^T \Sigma_{Y,\ln(Y!)} \boldsymbol{\nu} * Z \boldsymbol{\gamma}^{(m-1)} + X^T (Y - \mathbb{E}[Y]) \quad (5) \end{aligned}$$

and

$$
\begin{aligned}
-\boldsymbol{\nu} * Z^T \Sigma_{Y,\ln(Y!)} X\boldsymbol{\beta}^{(m)} + \boldsymbol{\nu}^2 * Z^T \Sigma_{\ln(Y!)} Z\boldsymbol{\gamma}^{(m)} \;=\; & -\boldsymbol{\nu} Z^T \Sigma_{Y,\ln(Y!)} X\boldsymbol{\beta}^{(m-1)} + \\
& \boldsymbol{\nu}^2 * Z^T \Sigma_{\ln(Y!)} Z\boldsymbol{\gamma}^{(m-1)} + \\
& \boldsymbol{\nu} * Z^T(-\ln(Y!) + \mathbb{E}[\ln(Y!)]) \quad (6)
\end{aligned}
$$

Each of the two equations in (5) and (6) are complicated and contain updates for both parameters. However, if we fix one parameter at each equation, a nice closed form expression appears for the updates. When we fix $\boldsymbol{\gamma}$ in equation (5) the equation reduces to

$$X^T \Sigma_Y X\boldsymbol{\beta}^{(m)} = X^T \Sigma_Y X\boldsymbol{\beta}^{(m-1)} + X^T(Y - \mathbb{E}[Y]). \quad (7)$$

This equation is nothing but WLS of X on Y with weights Σ_Y. Similarly, if we fix $\boldsymbol{\beta}$ in equation (6) then the equation reduces to

$$\boldsymbol{\nu}^2 * Z^T \Sigma_{\ln(Y!)} Z\boldsymbol{\gamma}^{(m)} = \boldsymbol{\nu}^2 * Z^T \Sigma_{\ln(Y!)} Z\boldsymbol{\gamma}^{(m-1)} + \boldsymbol{\nu} * Z^T(-\ln(Y!) + \mathbb{E}[\ln(Y!)]). \quad (8)$$

Again this is a WLS of $\boldsymbol{\nu} * Z$ on $\ln(Y!)$ with weights $\Sigma_{\ln(Y!)}$.

The two equations (7) and (8) are elegant and can be easily estimated with WLS methods. This approach is not only convenient for estimation but also helpful for generalizing to other modeling extensions such as additive models and the lasso.

3.3 Proof of Convergence of the Two-Step Method

To prove the convergence properties of our proposed two step algorithm we start with the following assumptions. Consider the parameter space $\Theta \in (0, \infty) \times (0, \infty)$, and the likelihood function L.

(A1). Let $\hat{\boldsymbol{\theta}}_{\mathbf{0}} = (\hat{\boldsymbol{\lambda}}_{\mathbf{0}}, \hat{\boldsymbol{\nu}}_{\mathbf{0}})^T \in \Theta$ be a starting value, such that $D_0 = \{\boldsymbol{\theta} = (\boldsymbol{\lambda}, \boldsymbol{\nu})^T \in \Theta | L(\boldsymbol{\theta}) \geq L(\hat{\boldsymbol{\theta}}_{\mathbf{0}})\}$ is compact.

(A2). The function L is uniquely maximized over D_0 for $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$.

(A3). Suppose that we have given parameter functions $\psi_i : D_0 \to \Theta_i (i = 1, 2)$ and let $M_i(\boldsymbol{\theta}), \boldsymbol{\theta} \in D_0$ be the corresponding sections: $M_i(\boldsymbol{\theta}) = \{\boldsymbol{\eta} \in D_0 | \psi_i(\boldsymbol{\eta}) = \psi_i(\boldsymbol{\theta})\} (i = 1, 2)$ Then we assume that, for $i = 1, 2$ and $\boldsymbol{\theta} \in D_0$, L is maximized uniquely by $T_i(\boldsymbol{\theta})$ on the section $M_i(\boldsymbol{\theta})$ and that $T_i(\boldsymbol{\theta})$ is continuous on D_0.

(A4). The point of global maximum $\hat{\boldsymbol{\theta}}$ is uniquely determined by the condition that it is the partial maximum along each section $M_i(\boldsymbol{\theta})$.

$$\sup_{\boldsymbol{\eta} \in M_i(\boldsymbol{\theta})} L(\boldsymbol{\eta}) = L(\boldsymbol{\theta}), i = 1, 2$$

implies $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$ or equivalently $T_i(\boldsymbol{\theta}) = \boldsymbol{\theta}$ implies $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$.

Assumptions A1 and A2 are based on the fact that the CMP distribution is unimodal and it has a log-concave pmf (Gupta et al. 2014, Steutel 1985). The remaining assumptions A3 and A4 follow from the properties of exponential family distributions. It is well known that the marginal distributions in a regular k-variate exponential family also belong to an exponential family (Keener 2006, Lehmann et al. 1991). It means that in an exponential family like CMP, the estimates obtained from maximizing the marginal likelihood are the same as the estimates obtained from maximizing the full likelihood.

Theorem 1 *Under assumptions A1-A4, the two step IRLS algorithm*

$$\hat{\boldsymbol{\theta}}_{\boldsymbol{n+1}} = T_1(T_2(\hat{\boldsymbol{\theta}}_{\boldsymbol{n}}))$$

converges to $\hat{\boldsymbol{\theta}}$ for any starting value in D_0.

The proof is similar to Jensen et al. (1991). The authors showed that under the above assumptions any partial maximization algorithm converges to true value for a given starting value.

Algorithm 1 IRLS Framework for CMP

Set initial values for $\boldsymbol{\nu}^{(0)}$ and $\boldsymbol{\lambda}^{(0)} = (Y + 0.1)^{\boldsymbol{\nu}^{(0)}}$.

Compute $\boldsymbol{\eta}_1^{(0)} = \ln(\boldsymbol{\lambda}^{(0)})$ and $\boldsymbol{\eta}_2^{(0)} = \ln(\boldsymbol{\nu}^{(0)})$

Compute $D^{(0)}(\boldsymbol{\lambda}^{(0)}, \boldsymbol{\nu}^{(0)}) = -2\ell(\boldsymbol{\lambda}^{(0)}, \boldsymbol{\nu}^{(0)})$.

Compute $\mathbb{E}[Y]^{(0)}$ and $\mathbb{V}[Y]^{(0)}$.

for i in 1:maxIter **do**

    Adjusted dependent variable: $T_1^{(i)} = \boldsymbol{\eta_1}^{(i-1)} + \frac{Y - \mathbb{E}[Y]^{(i-1)}}{\mathbb{V}[Y]^{(i-1)}}$

    Perform a weighted least squares regression of $T_1^{(i)}$ on X with weights $W_1 = \mathbb{V}(y)^{(i-1)}$ to obtain $\boldsymbol{\beta}^{(i)}$

    update $\boldsymbol{\eta_1}^{(i)} = X\boldsymbol{\beta}^{(i)}$ and $\boldsymbol{\lambda}^{(i)} = \exp(\boldsymbol{\eta_1}^{(i)})$

    Compute $\mathbb{E}[\ln(Y!)]^{(i-1)}$ and $\mathbb{V}[\ln(Y!)]^{(i-1)}$

    Adjusted dependent variable: $T_2^{(i)} = \boldsymbol{\nu}^{(i-1)} * \boldsymbol{\eta_2}^{(i-1)} + \frac{-\ln(Y!) + \mathbb{E}[\ln(Y!)]^{(i-1)}}{\mathbb{V}[\ln(Y!)]^{(i-1)}} * \boldsymbol{\nu}^{(i-1)}$

    Perform a weighted least squares regression of $T_2^{(i)}$ on $\boldsymbol{\nu}^{(i-1)} * Z$ with weights $W_2 = \mathbb{V}[\ln(Y!)]^{(i-1)}$ to obtain $\boldsymbol{\gamma}^{(i)}$

    update $\boldsymbol{\eta}_2^{(i)} = Z\boldsymbol{\gamma}^{(i)}$ and $\boldsymbol{\nu}^{(i)} = \exp(\boldsymbol{\eta}_2^{(i)})$

    Compute $D^{(i)}(\boldsymbol{\lambda}^{(i)}, \boldsymbol{\nu}^{(i)}) = -2\ell(\boldsymbol{\lambda}^{(i)}, \boldsymbol{\nu}^{(i)})$.

    if $\frac{D^{(i)} - D^{(i-1)}}{D^{(i)}} > 10^{-6}$ **then**

        Initiate step size optimization

    end if

    if $|\frac{D^{(i)} - D^{(i-1)}}{D^{(i)}}| < 10^{-6}$ **then**

        Convergence achieved. Break the loop.

    else

        Compute $\mathbb{E}[Y]^{(i)}$ and $\mathbb{V}[Y]^{(i)}$.

    end if

end for

3.4 Practical issues

3.4.1 Initial Values

Unlike other nonlinear optimization algorithms, IRLS does not require initial values for the parameters $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ but it does require initial values for $\boldsymbol{\lambda}$ and $\boldsymbol{\nu}$. We can provide suitable initial values based on the approximate method of a moment estimator for $\boldsymbol{\lambda}$ such as $(Y + 0.1)^{\boldsymbol{\nu}^{(o)}}$. However, we do not have a closed form expression for $\boldsymbol{\nu}^{(o)}$. In practice we observed that starting close to zero (for e.g $\boldsymbol{\nu} = 0.2$ or 0.5) yields satisfactory results.

3.4.2 Stopping Criterion

The standard IRLS algorithm uses the deviance as a stopping criterion. If the absolute relative change in the deviance is below some tolerance threshold, the algorithm stops. In general, the deviance for the i^{th} observation is defined as:

$$D_i = -2(\ell(y_i; \hat{\lambda}_i, \hat{\nu}_i) - \ell(y_i; \hat{\lambda}_{i,sat}, \hat{\nu}_{i,sat})).$$

The estimates for both $\hat{\lambda}_{i,sat}, \hat{\nu}_{i,sat}$ depend on each other and we do not have closed forms especially for the estimate $\hat{\nu}_{i,sat}$. For this reason, we consider only the term $-2\sum \ell(y_i; \hat{\lambda}_i, \hat{\nu}_i)$ and use it as our stopping criterion. Since the likelihood for the saturated model is constant across all the iterations, ignoring the term $\ell(y_i; \hat{\lambda}_{i,sat}, \hat{\nu}_{i,sat})$ does not impact our stopping criterion. In addition, this function is monotonic and if the algorithm is converging, it will decrease with every iteration.

3.4.3 Step Size

It is common for IRLS to exhibit convergence problems (Marschner et al. 2011). To avoid non-convergence issues we used the step-halving approach as suggested by Marschner et al.

(2011). The algorithm invokes step-halving either at the boundary or if the deviance is increasing. This step-halving makes sure that the algorithm remains in the interior space which is required for convergence.

3.5 Inference

Theorem 2 *Under regular conditions, the maximum likelihood estimators* $\hat{\boldsymbol{\theta}} = (\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})^T$ *are consistent and asymptotically normal:*

$$\sqrt{n}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \sim N(0, I^{-1}(\boldsymbol{\theta})).$$

Since the algorithm estimates each parameter separately while keeping the other parameter fixed, it only provides the marginal information for the respective parameters. The conditional information matrices can be obtained by correcting the marginal information as following:

$$\sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}|\hat{\boldsymbol{\gamma}}) \sim N(0, (I_{11} - I_{12}I_{22}^{-1}I_{21})^{-1})$$
$$\sqrt{n}(\hat{\boldsymbol{\gamma}} - \boldsymbol{\gamma}|\hat{\boldsymbol{\beta}}) \sim N(0, (I_{22} - I_{21}I_{11}^{-1}I_{12})^{-1}).$$

We note that the estimates for both $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ are not independent and inferences on one parameter will be influenced by the other estimate. For most practical applications, inference on the parameter $\boldsymbol{\beta}$ is of primary interest and usually the parameter $\boldsymbol{\gamma}$ will be treated as a nuisance parameter.

4 Simulation Study

We conducted an extensive simulation study to evaluate the performance of our proposed IRLS algorithm in comparison to existing likelihood-based methods for estimating CMP regression model.

At present there are two R packages (*CMPRegression* by Sellers & Lotze (2011); *CompGLM* by Pollock (2014)) for fitting the CMP regression model. Both use general purpose optimization functions to maximize the likelihood function. While these two R packages are technically the same, they differ in terms of their implementations. From now on, we denote these packages as Opt_1 and Opt_2 respectively. While Opt_1 does not uses log link for the $\boldsymbol{\nu}$ model, Opt_2 does use a log link and allow the user to model the $\boldsymbol{\nu}$ parameter as well. Similarly, while Opt_1 just provides the log likelihood to the optimization function, Opt_2 also provides the gradients. More importantly, in terms of computational issues, the support functions for Opt_1 were implemented in R and the support functions for Opt_2 were implemented in C++ which makes it work much faster.

Both Opt_1 and Opt_2 methods have some limitations. One obvious limitation is their inability to handle larger counts. If there are any outliers in the data (such as one or two extremely large counts) both methods do not work and give errors related to problems with the Hessian computation. Table 2 provides more insights on this issue. We have also observed that, sometimes, these methods exhibit convergence issues even if there are no outliers in the data. Unfortunately, we could not find any clear pattern on when such problems arise.

In order to provide a clear comparison of our method with the two aforementioned likelihood-based methods, we carefully constructed a simulated data set without any large counts so that none of these methods face any convergence issues. We considered sample size $n = 500$ and chose 4 covariates to include in the model. The covariates are simulated from normal and uniform distributions and also allowing a mild correlation between one pair of covariates ($x_1 \sim U(0,1), x_2 \sim N(0,0.5)$, $x_3 \sim N(0,0.1)$ and $x_4 = 0.2x_3 + N(0,0.5)$).

We considered four different values for ν to capture over dispersion, equi dispersion and under dispersion. The true values for the regression coefficients and their estimated values

At least one count	$IRLS$	Opt_1	Opt_2
< 90	Yes	Yes	Yes
$[90, 200]$	Yes	Yes	No
≥ 200	Yes	No	No

Table 2: Estimation of CMP regression with $IRLS$, Opt_1 and Opt_2 in the presence of at least one outlier. Rows display different magnitudes of the outlier value. Yes=works properly, No=returns an error.

using 20 bootstrap replications are reported in Table 3. From the results, it can be observed that IRLS performs equally well to the existing likelihood-based methods, especially Opt_1. While the three methods are indistinguishable for the over-dispersion case ($\nu = 0.5$), we observe that there are some clear discrepancies for the under-dispersion case ($\nu = 2.5$ or 4). In particular, Opt_2 has some issues when there is under-dispersion in the data.

For a couple of models from our simulation study ($\nu = 0.5, 4$) we compared the computation times of $IRLS$ with both Opt_1 and Opt_2 with increasing sample sizes. It is well known that the convergence speed of the IRLS algorithm depends on its starting value. We therefore take a subsample of data and then run the algorithm to get an approximated value for ν and feed it as the initial value for the estimation using the full data. We call this $IRLS_2$ and use $IRLS_1$ to denote the original algorithm which always starts at $\nu = 0.2$.

The computation times for the four methods are shown in Figures 3 and 4. From the results, it can be observed that Opt_2 is superior to the both $IRLS_1$, $IRLS_2$ and Opt_1. Opt_1 is painfully slow and often takes many minutes where both $IRLS_2$ and Opt_2 would have taken a few seconds. As expected, the $IRLS_2$ algorithm performs much faster than the original $IRLS_1$ and at times it even works faster than Opt_2. It is also worth mentioning

	$\nu = 0.5$			
	θ	$\hat{\theta}_{IRLS}$ (sd)	$\hat{\theta}_{Opt_1}$ (sd)	$\hat{\theta}_{Opt_2}$ (sd)
β_0	**0.05**	**0.05**	**0.04**	**0.04**
		(0.07)	(0.07)	(0.07)
β_1	**0.5**	**0.52**	**0.52**	**0.52**
		(0.07)	(0.07)	(0.07)
β_2	**-0.5**	**-0.52**	**-0.52**	**-0.52**
		(0.06)	(0.06)	(0.06)
β_3	**0.25**	**0.31**	**0.31**	**0.31**
		(0.21)	(0.21)	(0.21)
β_4	**-0.25**	**-0.25**	**-0.25**	**-0.25**
		(0.06)	(0.06)	(0.06)
$log(\nu)$	**-0.69**	**-0.65**	**-0.67**	**-0.67**
		(0.09)	(0.09)	(0.09)

	$\nu = 1$			
	θ	$\hat{\theta}_{IRLS}$ (sd)	$\hat{\theta}_{Opt_1}$ (sd)	$\hat{\theta}_{Opt_2}$ (sd)
β_0	**0.5**	**0.46**	**0.50**	**0.48**
		(0.11)	(0.69)	(0.06)
β_1	**1**	**0.99**	**1.03**	**1.01**
		(0.11)	(0.09)	(0.08)
β_2	**-1**	**-0.98**	**-1.03**	**-1**
		(0.09)	(0.08)	(0.06)
β_3	**0.5**	**0.56**	**0.59**	**0.58**
		(0.23)	(0.24)	(0.25)
β_4	**-0.5**	**-0.49**	**-0.51**	**-0.50**
		(0.08)	(0.08)	(0.07)
$log(\nu)$	**0**	**-0.03**	**0.02**	**0**
		(0.10)	(0.06)	(0)

	$\nu = 2.5$			
	θ	$\hat{\theta}_{IRLS}$ (sd)	$\hat{\theta}_{Opt_1}$ (sd)	$\hat{\theta}_{Opt_2}$ (sd)
β_0	**1**	**1.02**	**1.02**	**0.73**
		(0.13)	(0.13)	(0.16)
β_1	**3**	**3.09**	**3.09**	**2.43**
		(0.22)	(0.23)	(0.23)
β_2	**-3**	**-3.10**	**-3.10**	**-2.42**
		(0.20)	(0.20)	(0.24)
β_3	**2**	**2.17**	**2.17**	**1.76**
		(0.39)	(0.39)	(0.37)
β_4	**-2**	**-2.06**	**-2.06**	**-1.61**
		(0.17)	(0.17)	(0.16)
$log(\nu)$	**0.91**	**0.95**	**0.94**	**0.69**
		(0.06)	(0.06)	(0.10)

	$\nu = 4$			
	θ	$\hat{\theta}_{IRLS}$ (sd)	$\hat{\theta}_{Opt_1}$ (sd)	$\hat{\theta}_{Opt_2}$ (sd)
β_0	**2**	**2.00**	**2.01**	**1.01**
		(0.17)	(0.17)	(0.17)
β_1	**3**	**3.06**	**3.07**	**1.85**
		(0.23)	(0.24)	(0.16)
β_2	**-3**	**-3.08**	**-3.09**	**-1.88**
		(0.20)	(0.20)	(0.19)
β_3	**4**	**4.28**	**4.30**	**2.71**
		(0.45)	(0.46)	(0.49)
β_4	**-4**	**-4.08**	**-4.10**	**-2.38**
		(0.23)	(0.23)	(0.23)
$log(\nu)$	**1.38**	**1.41**	**1.40**	**0.86**
		(0.05)	(0.05)	(0.09)

Table 3: Comparison of the estimated parameters from three methods $(\theta_{IRLS}, \theta_{Opt_1}, \theta_{Opt_2})$. Results are obtained using 20 bootstrap replications. θ denotes the true parameter values.

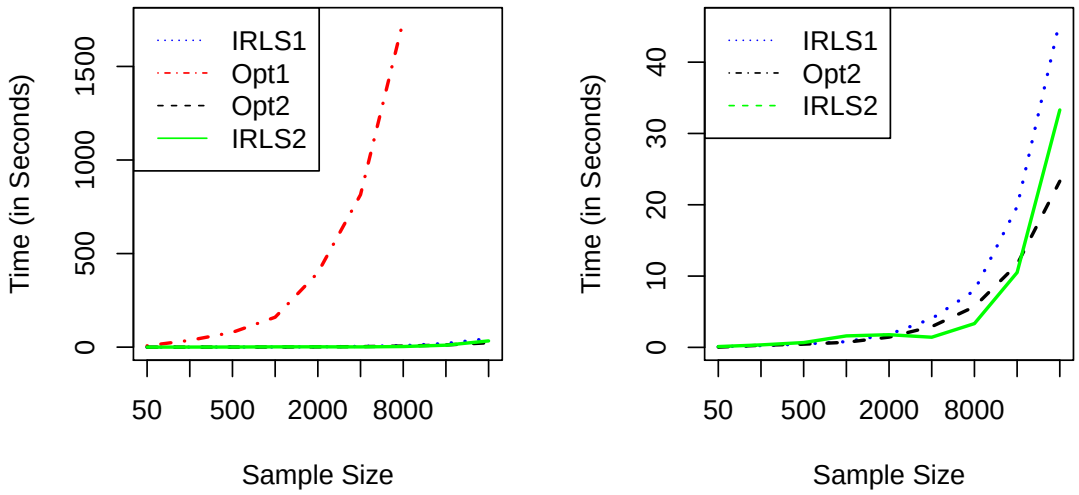

Figure 3: Comparison of the methods in terms of their computational timings for a data with $\nu = 0.5$ with increasing sample sizes. While $IRLS_1$ is initialized at $\nu = 0.2$, $IRLS_2$ is initialized with a ν computed from a sample model.

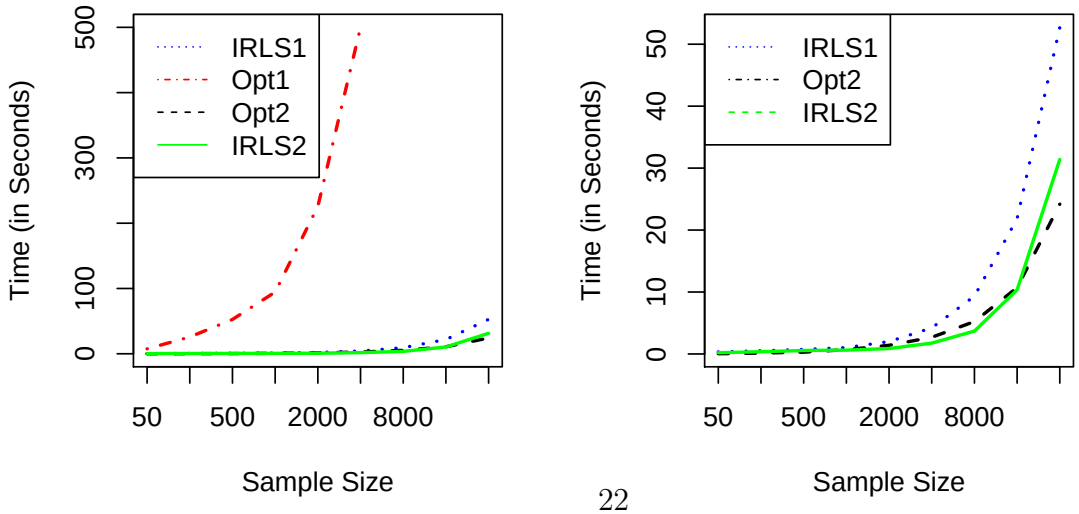

Figure 4: Comparison of the methods in terms of their computational timings for a data with $\nu = 4$ with increasing sample sizes. While $IRLS_1$ starts at $\nu = 0.2$, $IRLS_2$ starts with a ν computed from a sample model.

that while Opt_2 works very fast, it has some issues when there is under-dispersion as we have already seen in the simulation results.

5 A CMP Generalized Additive Model

Additive models are popular due to their flexibility for modeling nonlinear relationships while maintaining simplicity in interpretation. Similar to the existing GAMs for GLMs such as Poisson and logistic regression, the GAM for CMP can be defined as

$$\ln(\boldsymbol{\lambda}) = \boldsymbol{\alpha} + \sum_{j=1}^{p} \boldsymbol{f_j}(X_j)$$

$$\ln(\boldsymbol{\nu}) = Z\boldsymbol{\gamma}$$

where $\boldsymbol{f_j}$ $(j = 1, 2, \ldots, p)$ are the smooth functions for the p variables.

While there are multiple approaches for estimating $\boldsymbol{f_j}$ (Ruppert et al. 2003, Wood 2006, Hastie & Tibshirani 1990, 1986), we considered the approach of smoothing splines. Hastie & Tibshirani (1990) developed a common framework for estimating GAMs known as the *local scoring algorithm*. The algorithm can be used within the IRLS framework by incorporating another inner loop to estimate smooth terms at every iteration. Similar to other GLMs, it is straightforward to implement GAM for the CMP given an IRLS algorithm. Although the existing inference procedure for GAM is partly ad-hoc, we still use this inference procedure as a starting point to illustrate the usefulness of the additive model implementation.

Each step of the local scoring consists of a back-fitting loop applied to the adjusted dependent variable. The IRLS algorithm for CMP remains the same except that instead of WLS on $\boldsymbol{t_1}$ the local scoring is used to get the updated estimates of the smoothed functions.

In CMP, the adjusted dependent variable for the $\ln(\boldsymbol{\lambda})$ model is $\boldsymbol{t_1} = X\boldsymbol{\beta} + \Sigma_Y^{-1}(Y -$

$\mathbb{E}(Y)$). If R is the additive fit operator, then at the convergence

$$\begin{aligned}\hat{\boldsymbol{\eta}} = X\hat{\boldsymbol{\beta}} &= R(\hat{\boldsymbol{\eta}} + \Sigma_Y^{-1}(Y - \mathbb{E}(\hat{Y}))) \\ &= R\boldsymbol{t_1}.\end{aligned}$$

If the observations Y are i.i.d, then it is reasonable to assume that asymptotically $\boldsymbol{t_1}$ has mean $\boldsymbol{\eta}$ and variance $\Sigma_Y^{-1}\Sigma_Y\Sigma_Y^{-1}$.

We can write,

$$cov(\hat{\boldsymbol{\eta}}) \approx R\Sigma_Y^{-1}R^T$$
$$cov(\hat{\boldsymbol{f_j}}) \approx R_j\Sigma_Y^{-1}R_j^T$$

where R_j is the matrix that produces individual component smooth function $\hat{f}_j$ from $\boldsymbol{t_1}$. As mentioned in Chambers (1998), the standard error curves for nonparametric curve estimates are derived from the sum of two variances. One variance from the parametric part and another from the nonparametric part. Although this approach is not rigorous, its computational simplicity is appealing for use with CMP. Since the estimation of CMP itself is computationally heavier than common GLMs, we prefer to use a simpler approach such as this.

For inference on the smoothing terms, the existing *gam()* procedure in R Software (Hastie 2016) uses some type of score test approximations by calculating a change in a chi-square statistic by dropping each smooth term. We use this procedure to provide inference for the CMP smooth terms.

The back-fitting method is elegant as it has the flexibility to incorporate a wide variety of smoothing methods for component estimation. However, as suggested by Wood (2006), it is not easy to efficiently integrate the estimation of the smooth parameter into the model estimation framework. Common methods to estimate smooth parameters such as generalized cross-validation (GCV) might not work well especially with GLMs. While one could

estimate the global smooth matrix for the whole model, such estimation is computationally intensive. Wood (2006) proposed the penalized smoothing spline approach which can efficiently integrate the smoothness selection criterion into the estimation framework. Since we focus only on smoothing splines in this paper, we leave the topic of penalized splines with CMP for future work.

5.1 Simulation Study

We conducted a simulation study to evaluate the usefulness of the CMP GAM for fitting non linear terms. Inspired by the four uni-variate example Gu & Wahba (1991), Wood & Wood (2016), we simulated data from a CMP GAM, with sample size 500, as following:

- Simulate x_1, x_2, x_3 and x_4 from standard uniform $U[0, 1]$ distribution.
- Consider the functions $f_1 = sin(\pi x)$, $f_2 = \exp(x)$, $f_3 = 0.02x^2(1-x) + (0.5x)^2(1-x)^3$ and $f_4 = x$.
- Calculate $f = af_1(x_1) + bf_2(x_2) + cf_3(x_3)$, where a, b, c are pre-specified constants.
- Set $\lambda = \exp(f)$ and simulate data for a fixed ν.

To provide a comprehensive comparison we choose two different values of $\nu = \{0.5, 2.5\}$ which capture both over and under dispersion scenarios. While the constants a, b, c are set to $0.2, 0.5, -0.5$ respectively for $\nu = 0.5$, they are set to $1, 1, 1$ for $\nu = 2.5$.

We also considered a POISSON GAM and a CMP regression model for comparison with the CMP GAM. We used the $gam()$ function in R (Hastie 2016) to estimate the Poisson GAM. For each of the two scenarios ($\nu = \{0.5, 2.5\}$) we used 20 bootstrap replications and counted the significance levels for each nonparametric term in the model. We also recorded and compared their AIC. While the model equation for an additive model is

$y \sim s(x_1)+s(x_2)+s(x_3)+s(x_4)$, where $s(\cdot)$ is the smooth function, the model for the CMP regression is $y \sim x_1 + x_2 + x_3 + x_4$ without any smooth functions for covariates.

The simulation results are summarized in Table 4. The top table (a) describes the results for $\nu = 0.5$ and the bottom table (b) describes the results for $\nu = 2.5$. For the Table (a), from the AIC values it can be seen that the CMP GAM has better fit than both the Poisson GAM and the CMP regression. Comparing the significance levels for both additive models, the Poisson GAM exhibits higher levels of significance than the CMP GAM. This indicates that the Poisson GAM is not able to control for the over dispersion in the data. Ideally, the CMP regression would not produce significance levels for any of the coefficients because of the nonlinear terms in the true model. However, it can be observed that x_2 is significant and this is because the function $f_2 = \exp(x_2)$ is approximately equal to $1 + x_2$ (because $x_2 \in [0, 1]$). It is also worth noting that while x_4 is not part of the true model, the Poisson GAM identified it as significant while it is insignificant in the other two models.

Similarly, we can interpret the results from Table 4 (b). In terms of coefficient significance, the Poisson GAM is not able to capture the underlying true effect. For example, while $s(x_2)$ is significant in the CMP GAM at least half of the time, the Poisson GAM fails to capture the effect. In terms of the overall fit the CMP GAM again outperforms the two alternative models.

	#Significant (of 20 bootstraps) $(\leq 0.001, \leq 0.01, \leq 0.5, \leq 0.1, n.s)$		
	$cmp - gam$	$poisson - gam$	$cmp - regression$
$s(x_1)$	(**20**,0,0,0,0)	(**20**,0,0,0,0)	(0,0,1,**19**)
$s(x_2)$	(**8**,5,4,1,2)	(**17**,2,0,1,0)	(**20**,0,0,0,0)
$s(x_3)$	(0,1,0,4,**15**)	(1,4,1,**8**,6)	(0,0,0,1,**19**)
$s(x_4)$	(0,1,0,0,**19**)	(1,1,3,3,**12**)	(0,1,2,1,**16**)
	Estimation and Fit		
	$cmp - gam$	$poisson - gam$	$cmp - regression$
ν	**-0.63**(0.07)		**-0.70**(0.08)
AIC	**2714.78**(41.80)	**2840.12**(71.40)	**2745.96**(40.98)

(a) For $\nu = 0.5$

	#Significant (of 20 boostraps) $(\leq 0.001, \leq 0.01, \leq 0.5, \leq 0.1, n.s)$		
	$cmp - gam$	$poisson - gam$	$cmp - regression$
$s(x_1)$	(**20**,0,0,0,0)	(**17**,3,0,0,0)	(0,0,1,**19**)
$s(x_2)$	(1,2,4,4,**9**)	(0,0,1,0,**19**)	(**20**,0,0,0,0)
$s(x_3)$	(0,0,1,0,**19**)	(0,0,0,0,**20**)	(0,0,0,0,**20**)
$s(x_4)$	(0,0,1,0,**19**)	(0,0,0,0,**20**)	(0,0,0,3,**17**)
	Estimation and Fit		
	$cmp - gam$	$poisson - gam$	$cmp - regression$
ν	**0.88**(0.06)		**0.77**(0.05)
AIC	**1446.56**(36.55)	**1579.77**(18.69)	**1484.87**(32.60)

(b) For $\nu = 2.5$.

Table 4: Comparison of coefficient significance level and fit among CMP GAM, Poisson GAM, and CMP regression with 20 bootstrap replications. The model for gams is $y \sim s(x_1) + s(x_2) + s(x_3) + s(x_4)$ and for the regression $y \sim x_1 + x_2 + x_3 + x_4$

6 Bike sharing Data

We return to the bike sharing application described in Section 2, using the hourly data. We considered the following model for both counts of casual and registered users:

$$\begin{aligned}\ln(\boldsymbol{\lambda}) &= \beta_0 + \beta_1 hour + \beta_2 holiday + \beta_3 weekday + \beta_4 weathersit + \\ &\quad s(atemp) + s(hum) + s(windspeed) + s(day) \\ \ln(\boldsymbol{\nu}) &= \gamma_0.\end{aligned}$$

Since the attributes *atemp* and *temp* are highly correlated, we included only *atemp* in the model. We kept control variables such as *holiday* and *weathersit* as parametric terms and included other continuous variables of interest such as *hum* and *windspeed* as nonparametric components.

For this study we have only considered the January 2012 data. The same analysis can be repeated for every month or for every season. Since we have two outcome variables of interest we fit two models to this data. The first model is for the number of hourly rentals for registered users and the second model is for the number of hourly rentals for casual users. For the sake of comparison we also fit a Poisson GAM. The coefficient significance results are described in Tables 5 and 6. For brevity we did not include the coefficient significance results for the parametric components as they are all significant and similar across both the CMP and Poisson additive models. However, for the nonparametric components, the results are different for the CMP and Poisson GAMs. From Table 5 (model for registered users), it can be observed that while the Poisson GAM declared all the variables as significant, the CMP GAM only considered *day* as significant. Since the data is highly over dispersed, as can be observed from the large negative value for γ_0 (< 0 over dispersion; > 0 under dispersion), the CMP GAM fit is better than the Poisson fit. The AIC value for CMP

	$cmp - gam$ (Npar Chisq)	$poisson - gam$ (Npar Chisq)
(Intercept)		
s(day)	14.23^{**}	335.49^{***}
s(atemp)	7.55	158.01^{***}
s(hum)	7.62	183.55^{***}
s(windspeed)	6.60	148.61^{***}
γ_0	-3.25^{***}	
AIC	7419.05	19419.23
RMSE	50.27	50.83

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

Table 5: Comparison of CMP GAM and Poisson GAM in terms of coefficient significance and fit for # rentals for registered users in January 2012.

	$cmp - gam$ (Npar Chisq)	$poisson - gam$ (Npar Chisq)
(Intercept)		
s(day)	33.26^{***}	131.32^{***}
s(atemp)	6.34	18.92^{***}
s(hum)	16.43^{**}	66.16^{***}
s(windspeed)	8.27	38.38^{***}
γ_0	-1.58^{***}	
AIC	4064.52	5016.87
RMSE	7.17	7.56

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

Table 6: Comparison of CMP GAM and Poisson GAM in terms of coefficient significance and fit, for # rentals for casual users in January 2012.

GAM is almost one third of that from the Poisson GAM, thereby emphasizing the need for the more flexible CMP additive model.

Similar results are seen in the Table 6 which models the number of rentals for casual users. Again, in the Poisson GAM declared all the variables are significant, while in the CMP GAM *atemp* is insignificant. Similar to the registered users case, the data is highly over dispersed. These results therefore also highlight potential inference errors when fitting a Poisson additive model in the case of data with excessive dispersion.

To provide a more meaningful interpretation we use partial plots. The partial plots for

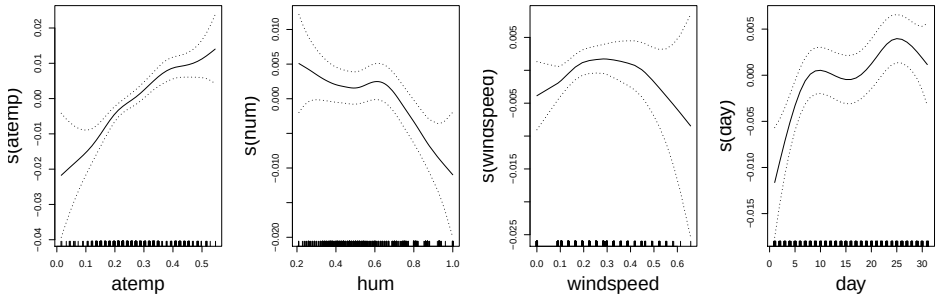

Figure 5: Partial plots for the smooth variables for CMP additive model. The dependent variable is # rentals for registered users in January 2012.

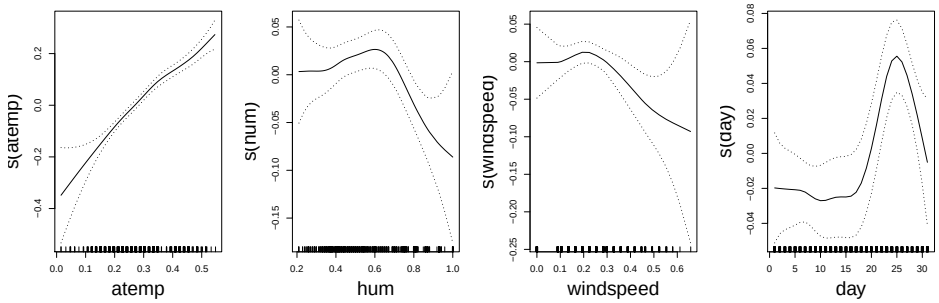

Figure 6: Partial plots for the smooth variables for CMP additive model. The dependent variable is #rentals for casual users in January 2012.

the CMP additive model are shown in Figures 5-6. For registered users, the smooth *day* is the only significant term among nonparametric smooth terms. It shows an increasing trend with a bump around January 15 and a decreasing trend at the month end.

For casual users one can draw interpretations from the partial plots in Figure 6. Unlike the registered users model, there are two smooth variables that are significant. The smooth variable *hum* show a hat shaped trend. There is an increasing trend until half-stage and decrease. The smooth variable *day* has an interesting trend. The trend is constant for the

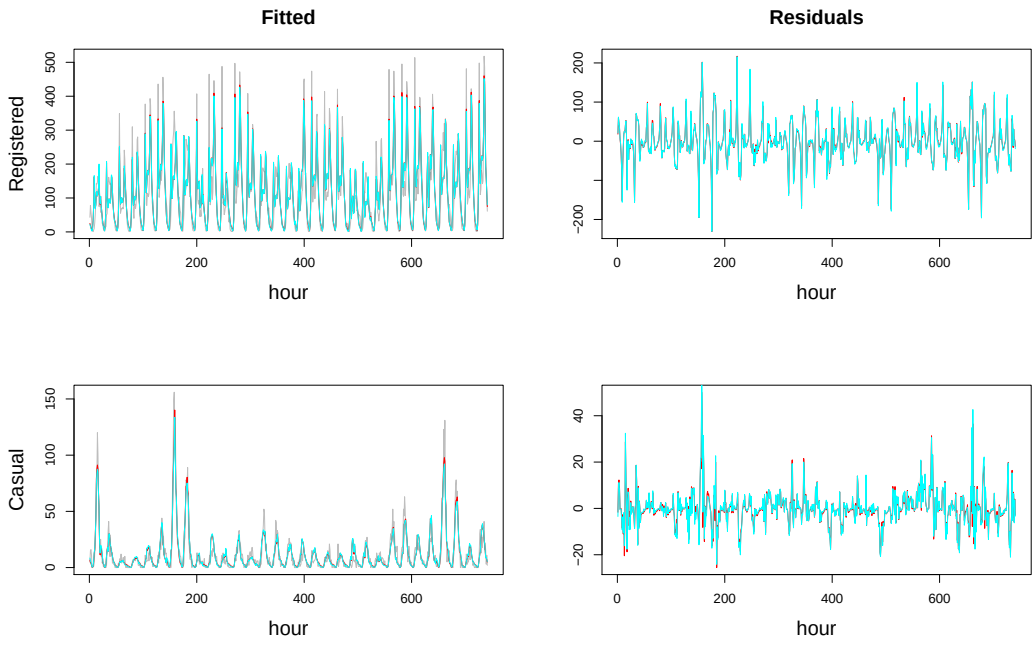

Figure 7: Comparison of the fitted values and residuals from CMP and Poisson additive models with actual values. The top plot is for registered users and the bottom plot is for casual users. (— true values, — CMP and —Poisson additive model).

first 15 days, followed by a sharp increase in rentals during days 20-25, and then again a decline. Understanding this behavior requires domain knowledge about bike sharing in Washington DC in January 2012.

Finally, from an actual model fit perspective, both CMP and Poisson additive models perform reasonably well. Figure 7 compares fitted values and residuals from both models. The RMSE values are given in Tables 5 and 6. We see that for both registered and casual user models the CMP and Poisson fitted values are close to the actual values. In fact, the CMP additive model slightly outperforms the Poisson additive model. This is also evident from the residual plots. The identical performance of the two models in terms

of fitted values is expected. However, if the prediction intervals are required, the results would be different from both models, as the CMP interval would account for the over- or under-dispersion by creating wider or narrower intervals, respectively Sellers & Shmueli (2010).

In summary, the CMP additive model can be a valuable addition for modeling count data. Despite its drawbacks, the CMP model is very flexible as it can handle both over dispersion and under dispersion which existing methods fail to handle. Although the bike sharing data did not exhibit under-dispersion, there are plenty of data sets that do. Moreover, when the researcher does not know the dispersion type (over or under) prior to modeling, CMP is a safe option.

7 Conclusions and Future Directions

We introduced a flexible estimation framework for estimating a CMP regression model that is based on the IRLS approach. This framework allows CMP to join other existing GLMs where IRLS is used for efficient estimation as well as for various modeling enhancements. This framework can be further developed to extend methods such as the Lasso.

While the IRLS algorithm for CMP is computationally intensive compared to an ordinary Poisson model, the computation time can be reduced by suitably parallelizing some of the computations such as the calculation of cumulants. Such parallel computing will be beneficial especially with large samples.

In this work we explored fitting additive models with smoothing splines. It is well known that the back fitting approach does not incorporate a smoothness selection criterion into the model estimation framework. It requires the user to manually fit different models for different values of the smoothing parameters. This process is computationally intensive

and time consuming especially with CMP. Therefore, there is need for a method which can estimate the smooth parameter along with fitting the model. One possible future direction is therefore developing an estimation framework with penalized splines which can integrate the smoothness selection criterion into the estimation process.

Finally, another open area for future research is predictive performance of the CMP GAM in terms of out-of-sample (future) prediction.

SUPPLEMENTARY MATERIAL

R-package We created an R-package (*cmp*) with all the methods developed in the paper. The package is available on github and can be installed by running the following R code:

```
require(devtools)
install_github("SuneelChatla/cmp")
require(cmp)
```

References

Abramowitz, M., Stegun, I. A. et al. (1966), 'Handbook of mathematical functions', *Applied mathematics series* **55**, 62.

Chambers, J. M. (1998), *Programming with data: A guide to the S language*, Springer Science & Business Media.

Dominici, F., McDermott, A., Zeger, S. L. & Samet, J. M. (2002), 'On the use of generalized additive models in time-series studies of air pollution and health', *American journal of epidemiology* **156**(3), 193–203.

Fanaee-T, H. & Gama, J. (2013), 'Event labeling combining ensemble detectors and background knowledge', *Progress in Artificial Intelligence* pp. 1–15.
URL: *http://dx.doi.org/10.1007/s13748-013-0040-3*

Gaunt, R. E., Iyengar, S., Daalhuis, A. B. & Simsek, B. (2016), 'An asymptotic expansion for the normalizing constant of the conway-maxwell-poisson distribution'. arXiv:1612.06618v1.

Green, P. J. (1984), 'Iteratively reweighted least squares for maximum likelihood estimation, and some robust and resistant alternatives', *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 149–192.

Gu, C. & Wahba, G. (1991), 'Minimizing gcv/gml scores with multiple smoothing parameters via the newton method', *SIAM Journal on Scientific and Statistical Computing* **12**(2), 383–398.

Gupta, R. C., Sim, S. Z. & Ong, S. H. (2014), 'Analysis of discrete data by conway–maxwell poisson distribution', *AStA Advances in Statistical Analysis* **98**(4), 327–343.
URL: *http://dx.doi.org/10.1007/s10182-014-0226-4*

Hastie, T. (2016), *gam: Generalized Additive Models*. R package version 1.14.
URL: *https://CRAN.R-project.org/package=gam*

Hastie, T. J. & Tibshirani, R. J. (1990), *Generalized additive models*, Vol. 43, CRC Press.

Hastie, T. & Tibshirani, R. (1986), 'Generalized additive models', *Statistical science* pp. 297–310.

Jensen, S. T., Johansen, S. & Lauritzen, S. L. (1991), 'Globally convergent algorithms for maximizing a likelihood function', *Biometrika* **78**(4), 867–877.

Keener, R. W. (2006), *Statistical theory: notes for a course in theoretical statistics*, Springer.

Lehmann, E. L., Casella, G. & Casella, G. (1991), *Theory of point estimation*, Wadsworth & Brooks/Cole Advanced Books & Software.

Marschner, I. C. et al. (2011), 'glm2: fitting generalized linear models with convergence problems', *The R journal* **3**(2), 12–15.

Pollock, J. (2014), *CompGLM: Conway-Maxwell-Poisson GLM and distribution functions*. R package version 1.0.
URL: *https://CRAN.R-project.org/package=CompGLM*

Ruppert, D., Wand, M. P. & Carroll, R. J. (2003), *Semiparametric regression*, number 12, Cambridge university press.

Sellers, K. F. & Shmueli, G. (2010), 'A flexible regression model for count data', *Annals of Applied Statistics* **4**(2), 943–961.

Sellers, K. & Lotze, T. (2011), 'Compoissonreg: Conway-maxwell poisson (com-poisson) regression', *R package version 0.3* **2**.

Shaheen, S., Guzman, S. & Zhang, H. (2010), 'Bikesharing in europe, the americas, and asia: past, present, and future', *Transportation Research Record: Journal of the Transportation Research Board* (2143), 159–167.

Shmueli, G., Minka, T. P., Kadane, J. B., Borle, S. & Boatwright, P. (2005), 'A useful distribution for fitting discrete data: revival of the conway–maxwell–poisson distribution', *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **54**(1), 127–142.

Simsek, B. & Iyengar, S. (2016), 'Approximating the conway-maxwell-poisson normalizing constant', *filomat* **30(4)**, 953–960.

Steutel, F. (1985), 'Log-concave and log-convex distributions', *Wiley StatsRef: Statistics Reference Online* .

Stieb, D. M., Judek, S. & Burnett, R. T. (2003), 'Meta-analysis of time-series studies of air pollution and mortality: update in relation to the use of generalized additive models', *Journal of the Air & Waste Management Association* **53**(3), 258–261.

Wood, S. (2006), *Generalized additive models: an introduction with R*, CRC press.

Wood, S. & Wood, M. S. (2016), 'Package mgcv', *R package version* pp. 1–7.

Yee, T. W. (2007), *Vector generalized linear and additive models*, Springer.